

Justifications in Constraint Handling Rules for Logical Retraction in Dynamic Algorithms

Thom Frühwirth

Ulm University, Germany
thom.fruehwirth@uni-ulm.de

Abstract. We present a straightforward source-to-source transformation that introduces justifications for user-defined constraints into the CHR programming language. Then a scheme of two rules suffices to allow for logical retraction (deletion, removal) of constraints during computation. Without the need to recompute from scratch, these rules remove not only the constraint but also undo all consequences of the rule applications that involved the constraint. We prove a confluence result concerning the rule scheme and show its correctness.

When algorithms are written in CHR, constraints represent both data and operations. CHR is already incremental by nature, i.e. constraints can be added at runtime. Logical retraction adds decrementality. Hence any algorithm written in CHR with justifications will become fully dynamic. Operations can be undone and data can be removed at any point in the computation without compromising the correctness of the result. We present two classical examples of dynamic algorithms, written in our prototype implementation of CHR with justifications that is available online: maintaining the minimum of a changing set of numbers and shortest paths in a graph whose edges change.

1 Introduction

Justifications have their origin in truth maintenance systems (TMS) [McA90] for automated reasoning. In this knowledge representation method, derived information (a formula) is explicitly stored and associated with the information it originates from by means of justifications. This dependency can be used to explain the reason for a conclusion (consequence) by its initial premises. With the help of justifications, conclusions can be withdrawn by retracting their premises. By this *logical retraction*, e.g. default reasoning can be supported and inconsistencies can be repaired by retracting one of the reasons for the inconsistency. An obvious application of justifications are dynamic constraint satisfaction problems (DCSP), in particular over-constrained ones [BM06].

In this work, we extend the applicability of logical retraction to arbitrary algorithms that are expressed in the programming language Constraint Handling Rules (CHR) [Frü09,Frü15,FR18]. To accomplish logical retraction, we have to be aware that CHR constraints can also be deleted by rule applications. These

constraints may have to be restored when a premise is retracted. With logical retraction, any algorithm written in CHR will become *fully dynamic*¹.

Minimum Example. Given a multiset of numbers $\text{min}(n_1), \text{min}(n_2), \dots, \text{min}(n_k)$. The constraint (predicate) $\text{min}(n_i)$ means that the number n_i is a candidate for the minimum value. The following CHR rule filters the candidates.

$$\text{min}(N) \ \backslash \ \text{min}(M) \ \langle \Rightarrow \ N < M \ | \ \text{true}.$$

The rule consists of a left-hand side, on which a pair of constraints has to be matched, a guard check $N < M$ that has to be satisfied, and an empty right-hand side denoted by **true**. In effect, the rule takes two **min** candidates and removes the one with the larger value (constraints after the \backslash symbol are deleted). Note that the **min** constraints behave both as operations (removing other constraints) and as data (being removed).

CHR rules are applied exhaustively. Here the rule keeps on going until only one, thus the smallest value, remains as single **min** constraint, denoting the current minimum. If another **min** constraint is added during the computation, it will eventually react with a previous *min* constraint, and the correct current minimum will be computed in the end. Thus the algorithm as implemented in CHR is incremental. It is not decremental, though: We cannot logically retract a *min* candidate. While removing a candidate that is larger than the minimum would be trivial, the retraction of the minimum itself requires to remember all deleted candidates and to find their minimum. With the help of justifications, this logical retraction will be possible automatically.

Contributions and Overview of the Paper. In the next section we recall syntax and operational semantics for CHR. Our contributions are as follows:

- We introduce CHR with justifications ($\text{CHR}^{\mathcal{J}}$) in Section 3. We enhance standard CHR programs with justifications by a source-to-source program transformation. We show the operational equivalence of rule applications in both settings. Thus $\text{CHR}^{\mathcal{J}}$ is a conservative extension of standard CHR.
- We define a scheme of two rules to enable logical retraction of constraints based on justifications in Section 4. We show that the rule scheme is confluent with each rule in any given program, independent of the confluence of that program. We prove correctness of logical retraction: the result of a computation with retraction is the same as if the constraint would never have been introduced in the computation.
- We present a proof-of-concept implementation of $\text{CHR}^{\mathcal{J}}$ in CHR and Prolog (available online) in Section 5. We discuss two classical examples for dynamic algorithms, maintaining the minimum of a changing set of numbers and maintaining shortest paths in a graph whose edges change.

The paper ends with discussion of related work in Section 6 and with conclusions and directions for future work.

¹ Dynamic algorithms for dynamic problems should not be confused with dynamic programming.

2 Preliminaries

We recall the abstract syntax and the equivalence-based abstract operational semantics of CHR in this section. Upper-case letters stand for (possibly empty) conjunctions of constraints in this paper.

2.1 Abstract Syntax of CHR

Constraints are relations, distinguished predicates of first-order predicate logic. We differentiate between two kinds of constraints: *built-in (pre-defined) constraints* and *user-defined (CHR) constraints* which are defined by the rules in a CHR program.

Definition 1. A *CHR program* is a finite set of rules. A (*generalized*) *simpagation rule* is of the form

$$r : H_1 \setminus H_2 \Leftrightarrow C|B$$

where r : is an optional *name* (a unique identifier) of a rule. In the rule *head* (left-hand side), H_1 and H_2 are conjunctions of user-defined constraints, the optional *guard* $C|$ is a conjunction of built-in constraints, and the *body* (right-hand side) B is a goal. A *goal* is a conjunction of built-in and user-defined constraints. A *state* is a goal. Conjunctions are understood as *multisets* of their conjuncts.

In the rule, H_1 are called the *kept constraints*, while H_2 are called the *removed constraints*. At least one of H_1 and H_2 must be non-empty. If H_1 is empty, the rule corresponds to a *simplification rule*, also written

$$s : H_2 \Leftrightarrow C|B.$$

If H_2 is empty, the rule corresponds to a *propagation rule*, also written

$$p : H_1 \Rightarrow C|B.$$

In this work, we restrict given CHR programs to rules without built-in constraints in the body except *true* and *false*. This restriction is necessary as long as built-in constraint solvers do not support the removal of built-in constraints.

2.2 Abstract Operational Semantics of CHR

Computations in CHR are sequences of rule applications. The operational semantics of CHR is given by the state transition system. It relies on a structural equivalence between states that abstracts away from technical details in a transition[RBF09,Bet14].

State equivalence treats built-in constraints semantically and user-defined constraints syntactically. Basically, two states are equivalent if their built-in constraints are logically equivalent (imply each other) and their user-defined constraints form syntactically equivalent multisets. For example,

$$X = <Y \wedge Y = <X \wedge c(X, Y) \equiv X = Y \wedge c(X, X) \not\equiv X = Y \wedge c(X, X) \wedge c(X, X).$$

For a state S , the notation S_{bi} denotes the built-in constraints of S and S_{ud} denotes the user-defined constraints of S .

Definition 2 (State Equivalence). Two states $S_1 = (S_{1bi} \wedge S_{1ud})$ and $S_2 = (S_{2bi} \wedge S_{2ud})$ are *equivalent*, written $S_1 \equiv S_2$, if and only if

$$\models \forall(S_{1bi} \rightarrow \exists \bar{y}((S_{1ud} = S_{2ud}) \wedge S_{2bi})) \wedge \forall(S_{2bi} \rightarrow \exists \bar{x}((S_{1ud} = S_{2ud}) \wedge S_{1bi}))$$

with \bar{x} those variables that only occur in S_1 and \bar{y} those variables that only occur in S_2 .

Using this state equivalence, the abstract CHR semantics is defined by a single transition (computation step). It defines the application of a rule. Note that CHR is a committed-choice language, i.e. there is no backtracking in the rule applications.

Definition 3 (Transition). Let the rule $(r : H_1 \setminus H_2 \Leftrightarrow C \mid B)$ be a variant² of a rule from a given program \mathcal{P} . The *transition (computation step)* $S \mapsto_r T$ is defined as follows, where S is called *source state* and T is called *target state*:

$$\frac{S \equiv (H_1 \wedge H_2 \wedge C \wedge G) \quad (r : H_1 \setminus H_2 \Leftrightarrow C \mid B) \in \mathcal{P} \quad (H_1 \wedge C \wedge B \wedge G) \equiv T}{S \mapsto_r T}$$

The goal G is called *context* of the rule application. It is left unchanged.

A *computation (derivation)* of a goal S in a program \mathcal{P} is a connected sequence $S_i \mapsto_{r_i} S_{i+1}$ beginning with the *initial state (query)* S_0 that is S and ending in a *final state (answer, result)* or the sequence is *non-terminating (diverging)*. We may drop the reference to the rules r_i to simplify the presentation. The notation \mapsto^* denotes the reflexive and transitive closure of \mapsto .

If the source state can be made equivalent to a state that contains the head constraints and the guard built-in constraints of a variant of a rule, then we delete the removed head constraints from the state and add the rule body constraints to it. Any state that is equivalent to this target state is in the transition relation.

The abstract semantics does not account for termination of inconsistent states and propagation rules. From a state with inconsistent built-in constraints, any transition is possible. If a state can fire a propagation rule once, it can do so again and again. This is called trivial non-termination of propagation rules.

Minimum Example, contd. Here is a possible transition from a state $S = (min(0) \wedge min(2) \wedge min(1))$ to a state $T = (min(0) \wedge min(1))$:

$$\frac{S \equiv (min(X) \wedge min(Y) \wedge X \leq Y \wedge (X = 0 \wedge Y = 2 \wedge min(1))) \quad (min(X) \setminus min(Y) \Leftrightarrow X \leq Y \mid true) \quad (min(X) \wedge X \leq Y \wedge true \wedge (X = 0 \wedge Y = 2 \wedge min(1))) \equiv T}{S \mapsto T}$$

² A variant (renaming) of an expression is obtained by uniformly replacing its variables by fresh variables.

3 CHR with Justifications (CHR^ℳ)

We present a conservative extension of CHR by justifications. If they are not used, programs behave as without them. Justifications annotate atomic CHR constraints. A simple source-to-source transformation extends the rules with justifications.

Definition 4 (CHR Constraints and Initial States with Justifications). A *justification* f is a unique identifier. Given an atomic CHR constraint G , a *CHR constraint with justifications* is of the form G^F , where F is a set of justifications. An *initial state with justifications* is of the form $\bigwedge_{i=1}^n G_i^{\{f_i\}}$ where the f_i are distinct justifications.

We now define a source-to-source translation from rules to rules with justifications. Let *kill* and *rem* (remove) be unary *reserved* CHR constraint symbols. This means they are only allowed to occur in rules as specified in the following.

Definition 5 (Translation to Rules with Justifications). Given a generalized simpagation rule

$$r : \bigwedge_{i=1}^l K_i \setminus \bigwedge_{j=1}^m R_j \Leftrightarrow C \mid \bigwedge_{k=1}^n B_k$$

Its translation to a *simpagation rule with justifications* is of the form

$$rf : \bigwedge_{i=1}^l K_i^{F_i} \setminus \bigwedge_{j=1}^m R_j^{F_j} \Leftrightarrow C \mid \bigwedge_{j=1}^m \text{rem}(R_j^{F_j})^F \wedge \bigwedge_{k=1}^n B_k^F \text{ where } F = \bigcup_{i=1}^l F_i \cup \bigcup_{j=1}^m F_j.$$

The translation ensures that the head and the body of a rule mention exactly the same justifications. More precisely, each CHR constraint in the body is annotated with the union of all justifications in the head of the rule, because its creation is caused by all of the head constraints. The constraint *rem/1* (remember removed) stores the constraints removed by the rule together with their justifications.

3.1 Operational Equivalence of Rule Applications

Let $A, B, C \dots$ be states. For convenience, we will often consider them as multisets of atomic constraints. Then the notation $A - B$ denotes multiset difference, A without B . To avoid clutter, let $A^{\mathcal{J}}, B^{\mathcal{J}}, C^{\mathcal{J}} \dots$ stand for conjunctions (or corresponding states) whose atomic CHR constraints are annotated with justifications according to the above definition of the rule scheme. Similarly, let $\text{rem}(R)^{\mathcal{J}}$ denote a conjunction $\bigwedge_{j=1}^m \text{rem}(R_j^{F_j})^F$.

We show that rule applications correspond to each other in standard CHR and in CHR^ℳ.

Lemma 1 (Equivalence of Program Rules). There is a computation step $S \mapsto_r T$ with simpagation rule

$$r : H_1 \setminus H_2 \Leftrightarrow C | B$$

if and only if there is a computation step with justifications $S^{\mathcal{J}} \mapsto_{rf} T^{\mathcal{J}} \wedge \text{rem}(H_2)^{\mathcal{J}}$ with the corresponding simpagation rule with justifications

$$rf : H_1^{\mathcal{J}} \setminus H_2^{\mathcal{J}} \Leftrightarrow C | \text{rem}(H_2)^{\mathcal{J}} \wedge B^{\mathcal{J}}.$$

Proof. We compare the two transitions involving rule r and rf , respectively:

$$\frac{(r : H_1 \setminus H_2 \Leftrightarrow C | B) \quad S \equiv (H_1 \wedge H_2 \wedge C \wedge G) \quad (H_1 \wedge C \wedge B \wedge G) \equiv T}{S \mapsto_r T}$$

$$\frac{(rf : H_1^{\mathcal{J}} \setminus H_2^{\mathcal{J}} \Leftrightarrow C | \text{rem}(H_2)^{\mathcal{J}} \wedge B^{\mathcal{J}}) \quad S^{\mathcal{J}} \equiv (H_1^{\mathcal{J}} \wedge H_2^{\mathcal{J}} \wedge C \wedge G^{\mathcal{J}}) \quad (H_1^{\mathcal{J}} \wedge C \wedge B^{\mathcal{J}} \wedge G^{\mathcal{J}}) \equiv T^{\mathcal{J}} \wedge \text{rem}(H_2)^{\mathcal{J}}}{S^{\mathcal{J}} \mapsto_{rf} T^{\mathcal{J}} \wedge \text{rem}(H_2)^{\mathcal{J}}}$$

Given the standard transition with rule r , the corresponding transition with justifications with rule rf is always possible. The rule rf by definition involves the same constraints (up to addition of justifications) as the standard rule r . It does not impose any additional built-in constraints, in particular no constraints on its justifications. The justifications in the rule body are computed as the union of the justifications in the rule head, which is always possible. Furthermore, the rem constraints in the rule body are simply added to the target state.

Conversely, given the transition with justifications with rule rf , we can strip away³ all justifications from it and remove $\text{rem}(H_2)^{\mathcal{J}}$ from the rule and the target state to arrive at the standard transition with rule r . \square

Since computations are sequences of connected computation steps, this lemma implies that computations for program rules in CHR and in CHR with justifications correspond to each other. Thus $\text{CHR}^{\mathcal{J}}$ is a conservative extension of CHR.

4 Logical Retraction Using Justifications

We use justifications to remove a CHR constraint from a computation without the need to recompute from scratch. This means that all its consequences due to rule applications it was involved in are undone. CHR constraints added by those rules are removed and CHR constraints removed by the rules are re-added. To specify and implement this behavior, we give a scheme of two rules, one for retraction and one for re-adding of constraints. The reserved CHR constraint $\text{kill}(f)$ undoes all consequences of the constraint with justification f .

³ For a related strip function, see the proof of Theorem 3.

Definition 6 (Rules for CHR Logical Retraction). For each n -ary CHR constraint symbol c (except the reserved *kill* and *rem*), we add a rule to kill constraints and a rule to revive removed constraints of the form:

$$\text{kill} : \text{kill}(f) \setminus G^F \Leftrightarrow f \in F \mid \text{true}$$

$$\text{revive} : \text{kill}(f) \setminus \text{rem}(G^{F_c})^F \Leftrightarrow f \in F \mid G^{F_c},$$

where $G = c(X_1, \dots, X_n)$, where X_1, \dots, X_n are different variables.

Note that a constraint may be revived and subsequently killed. This is the case when both F_c and F contain the justification f .

4.1 Confluence of Logical Retraction

Confluence of a program guarantees that any computation starting from a given initial state can always reach equivalent states, no matter which of the applicable rules are applied. There is a decidable, sufficient and necessary syntactic condition to check confluence of terminating programs and to detect rule pairs that lead to non-confluence when applied.

Definition 7 (Confluence). If $A \mapsto^* B$ and $A \mapsto^* C$ then there exist states D_1 and D_2 such that $B \mapsto^* D_1$ and $C \mapsto^* D_2$ where $D_1 \equiv D_2$.

Theorem 1. [Abd97,AFM99] A terminating CHR program is confluent if and only if all its critical pairs are joinable.

Decidability comes from the fact that there is only a finite number of critical pairs to consider.

Definition 8 (Overlap, Critical Pair). Given two (not necessarily different) simpagation rules whose variables have been renamed apart, $K_1 \setminus R_1 \Leftrightarrow C_1 \mid B_1$ and $K_2 \setminus R_2 \Leftrightarrow C_2 \mid B_2$. Let A_1 and A_2 be non-empty conjunctions of constraints taken from $K_1 \wedge R_1$ and $K_2 \wedge R_2$, respectively. An *overlap* of the two rules is the state consisting of the rules heads and guards:

$$((K_1 \wedge R_1) - A_1) \wedge K_2 \wedge R_2 \wedge A_1 = A_2 \wedge C_1 \wedge C_2.$$

The *critical pair* are the two states that come from applying the two rules to the overlap, where $E = (A_1 = A_2 \wedge C_1 \wedge C_2)$:

$$(((K_1 \wedge K_2 \wedge R_2) - A_2) \wedge B_1 \wedge E) \langle \rangle ((K_1 \wedge R_1 \wedge K_2) - A_1) \wedge B_2 \wedge E).$$

Note that the two states in the critical pair differ by $R_2 \wedge B_1$ and $R_1 \wedge B_2$.

A critical pair is *trivially joinable* if its built-in constraints are inconsistent or if both A_1 and A_2 do not contain removed constraints [AFM99].

We are ready to show the confluence of the *kill* and *revive* rules with each other and with each rule in any given program. It is not necessary that the given program is confluent. This means for any given program, the order between applying applicable rules from the program and retracting constraints can be freely interchanged. It does not matter for the result, if we kill a constraint first or if we apply a rule to it and kill it and its consequences later.

Theorem 2 (Confluence of Logical Retraction). Given a CHR program whose rules are translated to rules with justifications together with the *kill* and *revive* rules. We assume there is at most one $kill(f)$ constraint for each justification f in any state. Then all critical pairs between the *kill* and *revive* rules and any rule from the program with justifications are joinable.

Proof. There is only one overlap between the *kill* and *revive* rules,

$$\text{kill} : kill(f) \setminus G^F \Leftrightarrow f \in F \mid true$$

$$\text{revive} : kill(f) \setminus rem(G^{F_c})^F \Leftrightarrow f \in F \mid G^{F_c},$$

since G^F cannot have the reserved constraint symbol $rem/1$. The overlap is in the $kill(f)$ constraint. But since it is not removed by any rule, the resulting critical pair is trivially joinable.

By our assumption, the only overlap between two instances of the *kill* rule must have a single $kill(f)$ constraint. Again, since it is not removed, the resulting critical pair is trivially joinable. The same argument applies to the only overlap between two instances of the *revive* rule.

Since the head of a simpagation rule with justifications from the given program

$$rf : K^{\mathcal{J}} \setminus R^{\mathcal{J}} \Leftrightarrow C \mid rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}}$$

cannot contain reserved *kill* and *rem* constraints, these program rules cannot have an overlap with the *revive* rule.

But there are overlaps between program rules, say a rule rf , and the *kill* rule. They take the general form:

$$kill(f) \wedge K^{\mathcal{J}} \wedge R^{\mathcal{J}} \wedge G^F = A^F \wedge f \in F \wedge C,$$

where A^F occurs in $K^{\mathcal{J}} \wedge R^{\mathcal{J}}$. This leads to the critical pair

$$(kill(f) \wedge ((K^{\mathcal{J}} \wedge R^{\mathcal{J}}) - G^F) \wedge E) \langle \rangle (kill(f) \wedge K^{\mathcal{J}} \wedge rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}} \wedge E),$$

where $E = (G^F = A^F \wedge f \in F \wedge C)$. In the first state of the critical pair, the *kill* rule has been applied and in the second state the rule rf . Note that A^F is atomic since it is equated to G^F in E . Since G^F has been removed in the first state and $G^F = A^F$, rule rf is no longer applicable in that state.

We would like to join these two states. The joinability between a rule rf and the *kill* rule can be visualized by the diagram:

$$\begin{array}{ccc}
 & kill(f) \wedge K^{\mathcal{J}} \wedge R^{\mathcal{J}} \wedge E & \\
 \swarrow^{kill} & & \searrow^{rf} \\
 kill(f) \wedge ((K^{\mathcal{J}} \wedge R^{\mathcal{J}}) - G^F) \wedge E^* & \xleftarrow{revive^*, kill^*} & kill(f) \wedge K^{\mathcal{J}} \wedge rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}} \wedge E
 \end{array}$$

We now explain this joinability result. The states of the critical pair differ. In the first state we have the constraints $R^{\mathcal{J}}$ and have G^F removed from $K^{\mathcal{J}} \wedge R^{\mathcal{J}}$,

while in the second state we have the body constraints $rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}}$ of rule rf instead. Any constraint in $rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}}$ must include f as justification by definition, because f occurred in the head constraint A^F and E contains $f \in F$.

The goal $rem(R)^{\mathcal{J}}$ contains rem constraints for each removed constraint from $R^{\mathcal{J}}$. But then we can use $kill(f)$ with the *revive* rule to replace all rem constraints by the removed constraints, thus adding $R^{\mathcal{J}}$ back again. Furthermore, we can use $kill(f)$ with the *revive* rule to remove each constraint in $B^{\mathcal{J}}$, as each constraint in $B^{\mathcal{J}}$ contains the justification f . So $rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}}$ has been removed completely and $R^{\mathcal{J}}$ has been re-added.

The two states may still differ in the occurrence of G^F (which is A^F). In the first state, G^F was removed by the *kill* rule. Now if A^F (G^F) was in $R^{\mathcal{J}}$, it has been revived with $R^{\mathcal{J}}$. But then the *kill* rule is applicable and we can remove A^F again. In the second state, if A^F was in $R^{\mathcal{J}}$ it has been removed together with $R^{\mathcal{J}}$ by application of rule rf . Otherwise, A^F is still contained in $K^{\mathcal{J}}$. But then the *kill* rule is applicable to A^F and removes it from $K^{\mathcal{J}}$. Now A^F (G^F) does not occur in the second state either.

We thus have arrived at the first state of the critical pair. Therefore the critical pair is joinable. \square

This means that given a state, if there is a constraint to be retracted, we can either kill it immediately or still apply a rule to it and use the *kill* and *revive* rules afterwards to arrive at the same resulting state.

Note that the confluence between the *kill* and *revive* rules and any rule from the program is independent of the confluence of the rules in the given program.

4.2 Correctness of Logical Retraction

We prove correctness of logical retraction. The result of a computation with retraction is the same as if the retracted constraint would never have been introduced in the computation. We show that given a computation starting from an initial state with a $kill(f)$ constraint that ends in a state where the *kill* and *revive* rules are not applicable, i.e. these rules have been applied to exhaustion, then there is a corresponding computation without constraints that contain the justification f .

Theorem 3 (Correctness of Logical Retraction). Given a computation

$$A^{\mathcal{J}} \wedge G^{\{f\}} \wedge kill(f) \mapsto^* B^{\mathcal{J}} \wedge rem(R)^{\mathcal{J}} \wedge kill(f) \not\mapsto_{kill,revive},$$

where f does not occur in $A^{\mathcal{J}}$. Then there is a computation without $G^{\{f\}}$ and $kill(f)$

$$A^{\mathcal{J}} \mapsto^* B^{\mathcal{J}} \wedge rem(R)^{\mathcal{J}}.$$

Proof. We distinguish between transitions that involve the justification f or do not. A rule that applies to constraints that do not contain the justification f will produce constraints that do not contain the justification. A rule application

that involves at least one constraint with a justification f will only produce constraints that contain the justification f .

We now define a mapping from a computation with $G^{\{f\}}$ to a corresponding computation without $G^{\{f\}}$. The mapping essentially strips away constraints that contain the justification f except those that are remembered by *rem* constraints. In this way, the exhaustive application of the *revive* and *kill* rules *kill*(f) is mimicked.

$$\begin{aligned} strip(f, A^{\mathcal{J}} \wedge B^{\mathcal{J}}) &:= strip(f, A^{\mathcal{J}}) \wedge strip(f, B^{\mathcal{J}}) \\ strip(f, rem(G^{F_1})^{F_2}) &:= strip(f, G^{F_1}) \text{ if } f \in F_2 \\ strip(f, G^F) &:= true \text{ if } G \text{ is an atomic constraint except } rem/1 \text{ and } f \in F \\ strip(f, G^F) &:= G^F \text{ otherwise.} \end{aligned}$$

We extend the mapping from states to transitions. We keep the transitions except where the source and target state are equivalent, in that case we replace the transition \mapsto by an equivalence \equiv . This happens when a rule is applied that involves the justification f . The mapping is defined in such a way that in this case the source and target state are equivalent. Otherwise a rule that does not involve f has been applied. The mapping ensures in this case that all necessary constraints are in the source and target state, since it keeps all constraints that do not mention the justification f . For a computation step $C^{\mathcal{J}} \mapsto D^{\mathcal{J}}$ we define the mapping as:

$$\begin{aligned} strip(f, C^{\mathcal{J}} \mapsto_{rf} D^{\mathcal{J}}) &:= strip(f, C^{\mathcal{J}}) \equiv strip(f, D^{\mathcal{J}}) \text{ if rule } rf \text{ involves } f \\ strip(f, C^{\mathcal{J}} \mapsto_{rf} D^{\mathcal{J}}) &:= strip(f, C^{\mathcal{J}}) \mapsto_{rf} strip(f, D^{\mathcal{J}}) \text{ otherwise.} \end{aligned}$$

We next have to show is that the mapping results in correct state equivalences and transitions. If a rule is applied that does not involve justification f , then it is easy to see that the mapping $strip(f, \dots)$ leaves states and transitions unchanged.

Otherwise the transition is the application of a rule rf from the program, the rule *kill* or the rule *revive* where f is contained in the justifications. Let the context $E^{\mathcal{J}}$ be an arbitrary goal where $f \in \mathcal{J}$. Then we have to compute

$$\begin{aligned} strip(f, kill(f) \wedge G^F \wedge f \in F \wedge E^{\mathcal{J}} \mapsto_{kill} kill(f) \wedge E^{\mathcal{J}}) \\ strip(f, kill(f) \wedge rem(G^{F_c})^F \wedge f \in F \wedge E^{\mathcal{J}} \mapsto_{revive} kill(f) \wedge G^{F_c} \wedge E^{\mathcal{J}}) \\ strip(f, K^{\mathcal{J}} \wedge R^{\mathcal{J}} \wedge C \wedge E^{\mathcal{J}} \mapsto_{rf} K^{\mathcal{J}} \wedge rem(R)^{\mathcal{J}} \wedge B^{\mathcal{J}} \wedge C \wedge E^{\mathcal{J}}) \end{aligned}$$

and to show that equivalent states are produced in each case. The resulting states are

$$\begin{aligned} true \wedge true \wedge true \wedge E^{\mathcal{J}'} &\equiv true \wedge E^{\mathcal{J}'} \\ true \wedge G^{F_c} \wedge true \wedge E^{\mathcal{J}'} &\equiv true \wedge G^{F_c} \wedge E^{\mathcal{J}'} \text{ if } f \notin F_c \\ true \wedge true \wedge true \wedge E^{\mathcal{J}'} &\equiv true \wedge true \wedge E^{\mathcal{J}'} \text{ if } f \in F_c \\ K^{\mathcal{J}'} \wedge R^{\mathcal{J}'} \wedge C \wedge E^{\mathcal{J}'} &\equiv K^{\mathcal{J}'} \wedge R^{\mathcal{J}'} \wedge C \wedge E^{\mathcal{J}'} \text{ where } f \notin \mathcal{J}', \end{aligned}$$

where, given a goal A , the expression $A^{\mathcal{J}'}$ contains all constraints from $A^{\mathcal{J}}$ that do not contain the justification f .

In the end state of the given computation we know that the *revive* and *kill* rules have been applied to exhaustion. Therefore all $rem(G^{F_1})^{F_2}$ where F_2 contains f have been replaced by G^{F_1} by the *revive* rule. Therefore all standard constraints with justification f have been removed by the *kill* rule (including those revived), just as we do in the mapping $strip(f, \dots)$.

Therefore the end states are indeed equivalent except for the remaining *kill* constraint. \square

5 Implementation

As a proof-of-concept, we implement CHR with justifications (CHR ^{\mathcal{J}}) in SWI-Prolog using its CHR library. This prototype source-to-source transformation is available online at <http://pmx.informatik.uni-ulm.de/chr/translator/>. The translated programs can be run in Prolog or online systems like WebCHR.

Constraints with Justifications. CHR constraints annotated by a set of justifications are realized by a binary infix operator **##**, where the second argument is a list of justifications:

$C^{\{F_1, F_2, \dots\}}$ is realized as `C ## [F1, F2, ...]`.

For convenience, we add rules that add a new justification to a given constraint C . For each constraint symbol c with arity n there is a rule of the form `addjust @ c(X1, X2, ... Xn) <=> c(X1, X2, ... Xn) ## [_F]`.

where the arguments of $X1, X2, \dots, Xn$ are different variables.

Rules with Justifications. A CHR simplification rule with justifications is realized as follows:

$$rf : \bigwedge_{i=1}^l K_i^{F_i} \setminus \bigwedge_{j=1}^m R_j^{F_j} \Leftrightarrow C \mid \bigwedge_{j=1}^m rem(R_j^{F_j})^F \wedge \bigwedge_{k=1}^n B_k^F \text{ where } F = \bigcup_{i=1}^l F_i \cup \bigcup_{j=1}^m F_j$$

`rf @ K1 ## FK1, ... \ R1 ## FR1, ... <=> C |
union([FK1, ... FR1, ...], Fs), rem(R1##FR1) ## Fs, ... B1 ## Fs, ...`

where the auxiliary predicate `union/2` computes the ordered duplicate-free union of a list of lists⁴.

Rules *remove* and *revive*. Justifications are realized as *flags* that are initially unbound logical variables. This eases the generation of new unique justifications and their use in killing. Concretely, the reserved constraint $kill(f)$ is realized as built-in equality $F=x$, i.e. the justification variable gets bound. If $kill(f)$ occurred in the head of a *kill* or *revive* rule, it is moved to the guard as equality test $F==x$. Note that we rename rule *kill* to **remove** in the implementation.

`revive : kill(f) \ rem(C^{F_c})^F \Leftrightarrow f \in F \mid C^{F_c}`
`kill : kill(f) \ C^F \Leftrightarrow f \in F \mid true`

⁴ More precisely, a simplification rule is generated if there are no kept constraints and a propagation rule is generated if there are no removed constraints.

```

revive @ rem(C##FC) ## Fs <=> member(F,Fs),F==r | C ## FC.
remove @ C ## Fs <=> member(F,Fs),F==r | true.

```

Since rules are tried in program order in the CHR implementation, the constraint `C` in the second rule is not a reserved `rem/1` constraint when the rule is applicable. The check for set membership in the guards is expressed using the standard nondeterministic Prolog built-in predicate `member/2`.

Logical Retraction with `killc/1`. We extend the translation to allow for retraction of derived constraints. The constraint `killc(C)` logically retracts one occurrence of a constraint `C`. The two rules `killc` and `killr` try to find the constraint `C`. The `killr` rule applies in the case where constraint `C` has been removed and is therefore now present in a `rem` constraint. The associated justifications point to all initial constraints that were involved in producing the constraint `C`. For retracting the constraint, it is sufficient to remove one of its producers. This introduces a choice which is implemented by the `member` predicate.

```

killc @ killc(C), C ## Fs <=> member(F,Fs),F=r.
killr @ killc(C), rem(C ## FC) ## _Fs <=> member(F,FC),F=r.

```

Note that in the `killr` rule, we bind a justification `F` from `FC`, because `FC` contains the justifications of the producers of constraint `C`, while `Fs` also contains those that removed it by a rule application.

5.1 Examples

We discuss two classical examples for dynamic algorithms, maintaining the minimum of a changing set of numbers and shortest paths when edges change.

Dynamic Minimum. Translating the minimum rule to one with justifications results in:

```

min(A)##B \ min(C)##D <=> A<C | union([B,D],E), rem(min(C)##D)##E.

```

The following shows an example query and the resulting answer in SWI-Prolog:

```

?- min(1)##[A], min(0)##[B], min(2)##[C].
rem(min(1)##[A])##[A,B], rem(min(2)##[C])##[B,C], min(0)##[B].

```

The constraint `min(0)` remained. This means that 0 is the minimum. The constraints `min(1)` and `min(2)` have been removed and are now remembered. Both have been removed by the constraint with justification `B`, i.e. `min(0)`.

We now logically retract with `killc` the constraint `min(1)` at the end of the query. The `killr` rule applies and removes `rem(min(1)##[A])##[A,B]`.

```

killr @ killc(C), rem(C ## FC) ## _Fs <=> member(F,FC),F=r.

```

In the rule body, the justification `A` is bound to `r` – to no effect, since there are no other constraints with this justification:

```
?- min(1)##[A], min(0)##[B], min(2)##[C], killc(min(1)).
rem(min(2)##[C])##[B,C], min(0)##[B].
```

On the other hand, if we retract the minimum `min(0)`, the `killc` rule

```
killc @ killc(C), C ## Fs <=> member(F,Fs),F=r
```

applies. It removes `min(0)##[B]` and binds justification `B`. The two `rem` constraints for `min(1)` and `min(2)` have justification `B` as well, so these two constraints are re-introduced by applications of rule `revive`

```
revive @ rem(C##FC) ## Fs <=> member(F,Fs),F=r | C ## FC.
```

The minimum rule applies to these two revived constraints. Note that `min(2)` is now removed by `min(1)` (before it was `min(0)`). The result is the updated minimum, which of course is 1:

```
?- min(1)##[A], min(0)##[B], min(2)##[C], killc(min(0)).
rem(min(2)##[C])##[A,C], min(1)##[B].
```

Dynamic Shortest Path. Given a graph with directed arcs `e(X,Y)`, we compute the lengths of the shortest paths between all pairs of reachable nodes:

```
% keep shorter of two paths from X to Y
pp @ p(X,Y,L1) \ p(X,Y,L2) <=> L1=<L2 | true.
% edges have a path of unit length
e @ e(X,Y) ==> p(X,Y,1).
% extend path in front by an edge
ep @ e(X,Y), p(Y,Z,L) ==> L1:=L+1 | p(X,Z,L1).
```

The corresponding rules in the translated program are:

```
pp@p(A,B,C)##D \ p(A,B,E)##F <=> C=<E |
                                union([D,F],G), rem(p(A,B,E)##F)##G.
e @e(A,B)##C ==> true | union([C],D), p(A,B,1)##D.
ep@e(A,B)##C,p(B,D,E)##F ==> G is E+1 | union([C,F],H),p(A,D,G)##H.
```

Here is a sample query and its resulting answer.

```
?- e(a,b)##[A], e(b,c)##[B], e(a,c)##[C].
rem(p(a, c, 2)##[A, B])##[A,B,C],
p(a, b, 1)##[A], e(a, b)##[A],
p(b, c, 1)##[B], e(b, c)##[B],
p(a, c, 1)##[C], e(a, c)##[C].
```

We see that a path of length 2 has been removed by the constraint `e(a,c)##[C]`, which produced a shorter path of length one. We next kill this constraint `e(a,c)`.

```
?- e(a,b)##[A], e(b,c)##[B], e(a,c)##[C], kill(e(a,c)).
p(a, b, 1)##[A], e(a, b)##[A],
p(b, c, 1)##[B], e(b, c)##[B],
p(a, c, 2)##[A,B].
```

Its path $p(a,c,1)$ disappears and the removed path $p(a,c,2)$ is re-added. We can see that the justifications of a path contains are those from the edges in that path. The same happens if we logically retract $p(a,c,1)$ instead of $e(a,c)$.

What happens if we remove $p(a,c,2)$ from the initial query? The `killr` rule applies. Since the path has two justifications, there are two computations generated by the `member` predicate. In the first one, the constraint $e(a,b)$ disappeared, in the second answer, it is $e(b,c)$. In both cases, the path cannot be computed anymore, i.e. it has been logically retracted.

```
?- e(a,b)##[A], e(b,c)##[B], e(a,c)##[C], kill(p(a,c,2)).
p(b, c, 1)##[B], e(b, c)##[B],
p(a, c, 1)##[C], e(a, c)##[C]
;
p(a, b, 1)##[A], e(a, b)##[A],
p(a, c, 1)##[C], e(a, c)##[C].
```

6 Related Work

The idea of introducing justifications into CHR is not new. The thorough work of Armin Wolf on Adaptive CHR [WGG00] was the first to do so. Different to our work, this technically involved approach requires to store detailed information about the rule instances that have been applied in a derivation in order to undo them. In our approach, we use a straightforward source-to-source transformation and retract constraints one-by-one instead. Adaptive CHR had a low-level implementation in Java [Wol01], while we give an implementation in CHR itself by a straightforward source-to-source transformation that we prove confluent and correct. Moreover we prove confluence of the rule scheme for logical retraction with the rules of the given program. The application of adaptive CHR considered dynamic constraint satisfaction problems (DCSP) only, in particular for the implementation of search strategies [Wol05], while we apply our approach to arbitrary algorithms in order to make them fully dynamic.

The issue of search strategies was further investigated by Leslie De Koninck et. al. [DKSD08]. They introduce a flexible search framework in CHR^\vee (CHR with disjunction) extended with rule and search branch priorities. In their work, justifications are introduced into the semantics of CHR^\vee to enable dependency-directed backtracking in the form of conflict-directed backjumping. Our work does not need a new semantics for CHR, nor its extension with disjunction, it rather relies on a source-to-source transformation within the standard semantics.

Our work does not have a particular application of justifications in mind, but rather provides the basis for any type of application that requires dynamic algorithms. There are, however, CHR applications that use justifications. The work of Jeremy Wazny et. al. [SSW03] introduced informally a particular kind of justifications into CHR for the specific application of type debugging and reasoning in Haskell. Justifications correspond to program locations in the given Haskell program. Unlike other work, the constraints in the body of CHR rules have given justifications to which justifications from the rule applications are

added at runtime. The more recent work of Gregory Duck [Duc12] introduces SMCHR, a tight integration of CHR with a Boolean Satisfiability (SAT) solver for quantifier-free formulae including disjunction and negation as logical connectives. It is mentioned without giving further details that for clause generation, SMCHR supports justifications for constraints that include syntactic equality constraints between variables.

7 Conclusions

In this paper, the basic framework for CHR with justifications ($\text{CHR}^{\mathcal{J}}$) has been established and formally analyzed. We defined a straightforward source-to-source program transformation that introduces justifications into CHR as a conservative extension. Justifications enable logical retraction of constraints. If a constraint is retracted, the computation is adapted and continues as if the constraint never was introduced. We proved confluence and correctness of the two-rule scheme that encodes the logical retraction. We presented a prototype implementation that is available online together with two classical examples.

Our approach applies to CHR rules without built-in constraints in the body, because built-in constraint solvers typically do not support removal of constraints. But note that built-in constraints can be re-implemented in CHR. As one reviewer pointed out, CHR systems typically feature a propagation history to avoid the re-application of propagation rules and try to apply rules in the given order. This may lead to adverse effects in our implementation. The reviving of constraints may trigger redundant firings of propagation rules because the propagation history is partially lost then. The order in which constraints are revived may lead to different rule applications and thus to different results if the program is not confluent. We have not observed these effects in our experiments so far.

Future work could investigate three topics: dynamic algorithms, implementation and application domains of CHR with justifications. First, we would like to research how logical as well as classical algorithms implemented in CHR behave when they become dynamic. Second, we would like to improve the implementation, optimize and benchmark it. A first step can be found in the companion paper [Fru17] together with an informal complexity analysis. A formal complexity analysis would clarify the overhead involved in adding justifications. Currently, the entire history of removed constraints is stored. It could suffice to remember only a partial history if only certain constraints can be retracted or if partial recomputation proves to be more efficient for some constraints. A lower-level implementation could benefit from the propagation history of propagation rules. Third, the rule scheme can be extended to support typical application domains of justifications: explanation of derived constraints by justifications (for debugging), detection and repair of inconsistencies (for error diagnosis), and implementing nonmonotonic logical behaviors (e.g. default logic, abduction, defeasible reasoning).

Acknowledgements. We thank Daniel Gall for implementing the online transformation tool for $\text{CHR}^{\mathcal{J}}$. We also thank the anonymous reviewers for their helpful suggestions on how to improve the paper.

References

- [Abd97] Slim Abdennadher. Operational semantics and confluence of constraint propagation rules. In G. Smolka, editor, *CP '97: Proc. Third Intl. Conf. Principles and Practice of Constraint Programming*, volume 1330 of *LNCS*, pages 252–266. Springer, 1997.
- [AFM99] Slim Abdennadher, Thom Frühwirth, and Holger Meuss. Confluence and semantics of constraint simplification rules. *Constraints*, 4(2):133–165, 1999.
- [Bet14] Hariolf Betz. *A unified analytical foundation for constraint handling rules*. BoD, 2014.
- [BM06] Kenneth N Brown and Ian Miguel. Uncertainty and change, chapter 21. *Handbook of Constraint Programming*, pages 731–760, 2006.
- [DKSD08] Leslie De Koninck, Tom Schrijvers, and Bart Demoen. A flexible search framework for chr. In *Constraint Handling Rules — Current Research Topics*, volume LNAI 5388, pages 16–47. Springer, 2008.
- [Duc12] Gregory J Duck. SMCHR: Satisfiability modulo constraint handling rules. *Theory and Practice of Logic Programming*, 12(4-5):601–618, 2012.
- [FR18] Thom Frühwirth and Frank Raiser, editors. *Constraint Handling Rules - Compilation, Execution, and Analysis*. BOD, ISBN 9783746069050, January 2018.
- [Frü09] Thom Frühwirth. *Constraint Handling Rules*. Cambridge University Press, 2009.
- [Frü15] Thom Frühwirth. Constraint handling rules – what else? In *Rule Technologies: Foundations, Tools, and Applications*, pages 13–34. Springer International Publishing, 2015.
- [Fru17] Thom Frühwirth. Implementation of Logical Retraction in Constraint Handling Rules with Justifications. In *21st International Conference on Applications of Declarative Programming and Knowledge Management (INAP 2017)*, September 2017.
- [McA90] David A McAllester. Truth maintenance. In *AAAI*, volume 90, pages 1109–1116, 1990.
- [RBF09] Frank Raiser, Hariolf Betz, and Thom Frühwirth. Equivalence of CHR states revisited. In F. Raiser and J. Sneyers, editors, *CHR '09*, pages 33–48. K.U.Leuven, Dept. Comp. Sc., Technical report CW 555, July 2009.
- [SSW03] Peter J Stuckey, Martin Sulzmann, and Jeremy Wazny. Interactive type debugging in haskell. In *Proceedings of the 2003 ACM SIGPLAN workshop on Haskell*, pages 72–83. ACM, 2003.
- [WGG00] Armin Wolf, Thomas Gruenhagen, and Ulrich Geske. On the incremental adaptation of CHR derivations. *Applied Artificial Intelligence*, 14(4):389–416, 2000.
- [Wol01] Armin Wolf. Adaptive constraint handling with chr in java. In *International Conference on Principles and Practice of Constraint Programming*, pages 256–270. Springer, 2001.
- [Wol05] Armin Wolf. Intelligent search strategies based on adaptive constraint handling rules. *Theory and Practice of Logic Programming*, 5(4-5):567–594, 2005.